arXiv:1406.5706v2 [math.ST] 21 Sep 2014

On the Maximum Entropy Property of the First–Order Stable Spline Kernel and its Implications

Francesca P. Carli *
fpc23@cam.ac.uk.

Abstract

A new nonparametric approach for system identification has been recently proposed where the impulse response is seen as the realization of a zero–mean Gaussian process whose covariance, the so–called stable spline kernel, guarantees that the impulse response is almost surely stable. Maximum entropy properties of the stable spline kernel have been pointed out in the literature. In this paper we provide an independent proof that relies on the theory of matrix extension problems in the graphical model literature and leads to a closed form expression for the inverse of the first order stable spline kernel as well as to a new factorization in the form $UWU^\top$ with U upper triangular and W diagonal. Interestingly, all first–order stable spline kernels share the same factor U and W admits a closed form representation in terms of the kernel hyperparameter, making the factorization computationally inexpensive. Maximum likelihood properties of the stable spline kernel are also highlighted. These results can be applied both to improve the stability and to reduce the computational complexity associated with the computation of stable spline estimators.

1 Introduction

Most of the currently used techniques for linear system identification relies on parametric prediction error methods (PEMs), [Ljung, 1999, Soderstrom and Stoica, 1989]. Here, finite–dimensional hypothesis spaces of different order, such as ARX, ARMAX or Laguerre models, are first postulated. Then, the most adequate model order is selected trading–off between bias and variance to avoid overfitting. Model–order selection is usually performed by optimizing some penalized goodness–of–fit criteria, such as the Akaike information criterion (AIC) [Akaike, 1974] or the Bayesian information criterion (BIC) [Schwarz, 1978], or via cross validation (CV) [Hastie et al., 2008]. Statistical properties of prediction error methods are well understood under the assumption that the model class is fixed. Nevertheless, sample properties of PEM approaches equipped e.g. with AIC or CV can much depart from those predicted by standard (i.e. without model selection) statistical theory ([Pillonetto and De Nicolao, 2010, Pillonetto et al., 2011]).

Motivated by these pitfalls, a new approach to system identification has been recently proposed where the system impulse response is seen as the realization of a zero–mean

*Department of Electrical Engineering and Computer Science, University of Liège, Belgium and Department of Engineering, University of Cambridge, United Kingdom.

Gaussian process with a suitable covariance that depends on few hyperparameters, learnt from data via, e.g., marginal likelihood maximization. This procedure can be seen as the counterpart of model order selection in the parametric paradigm and in many cases it has been proved to be more robust than AIC-type criteria and CV.

In this scheme, quality of the estimates crucially depends on the covariance (kernel) of the Gaussian process. A large variety of positive semidefinite kernels have been introduced in the machine learning literature [Shawe-Taylor and Cristianini, 2004, Scholkopf and Smola, 2001]. Nevertheless, a straight application of standard machine learning kernels in the framework of system identification is doomed to fail mainly because of the lack of constraints on system stability. For this reason, several kernels have been recently introduced in the system identification literature [Pillonetto and De Nicolao, 2010, Chen et al., 2011].

This paper deals with *stable spline kernels*. Stable spline kernels were introduced in [Pillonetto and De Nicolao, 2010] as an adaptation of spline kernels that enforces the associated process realizations to be asymptotically stable. Some theoretical results that assess robustness of this class of kernels are described in [Aravkin et al., 2014, Carli et al., 2012a]. Efficient numerical implementations are discussed in [Carli et al., 2012b, Chen and Ljung, 2013]. In this paper we concentrate on *first–order stable spline kernels* (see [Pillonetto et al., 2010] and also [Chen et al., 2012], where this class of kernels has also been introduced by using a totally different, deterministic argument). Maximum entropy properties of first–order stable spline kernels have been pointed out in [Pillonetto and De Nicolao, 2011]. In this paper, we provide an alternative proof of the maximum entropy property by resorting to an independent, algebraic argument that connects to the theory of matrix completion and, in particular, of band extension problems in the graphical models literature [Dempster, 1972, Grone et al., 1984, Dym and Gohberg, 1981, Gohberg et al., 1993, Dahl et al., 2008]. This alternative approach leads to a closed form expression for the inverse of the first–order stable spline kernel. A factorization of the first–order stable spline kernel in the form $UWU^\top$ with U upper triangular and W diagonal is also provided. Interestingly, all first–order stable spline kernels share the same factor U and W admits a closed form representation in terms of the kernel hyperparameter, making the factorization inexpensive from a computational point of view. Moreover it can be proved that the first–order stable spline kernel maximizes the likelihood among all covariances that satisfy certain conditional independence constraints. The above mentioned properties can for example be used both to improve stability and reduce the computational burden of computational schemes for the evaluation of the stable spline estimator.

The paper is organized as follows. In Section 2 the problem is introduced and Gaussian process regression via first order stable–spline kernels is briefly reviewed. In Section 3 relevant theory of matrix completion problems is introduced. Section 4 contains our main results. Section 5 ends the paper.

Notation. Let $\mathcal{S}_n$ denote the vector space of symmetric matrices of order n. We write $A \succeq 0$ (resp. $A \succ 0$) to denote that A is positive semidefinite (resp. positive definite). Moreover, we denote by I_k the identity matrix of order k, by $\mathbb{1}_k$ the k–dimensional vector of all ones, and by 0_k the k–dimensional vector of all zeroes. The diagonal matrix of order k with diagonal elements $\{a_1, a_2, \ldots, a_k\}$ will be denoted by $\mathrm{diag}\{a_1, a_2, \ldots, a_k\}$. If A is a square matrix of order n, for index sets $\beta \subseteq \{1, \ldots, n\}$ and $\gamma \subseteq \{1, \ldots, n\}$, we denote the submatrix that lies in the rows of A indexed by β and the columns indexed by γ as $A(\beta, \gamma)$.

If $\gamma = \beta$, the submatrix $A(\beta, \gamma)$ is abbreviated $A(\beta)$.

2 Linear system identification via Gaussian Process Regression

2.1 Statement of the problem

We consider the measurement model

$$y_t = \sum_{k=1}^{\infty} f_k u_{t-k} + e_t \tag{1}$$

where $\{y_t\}$ denote the noisy output samples of a discrete–time linear dynamical system fed with a known input $\{u_t\}$. $f = \{f_t\}_{t=1}^{\infty}$ is the unknown impulse response and $\{e_t\}$ is white Gaussian noise with variance σ^2. Suppose that N measurements are available. We can collect these measurements in the N–dimensional column vector $y = [y_1, \ldots, y_N]^\top$. Let e denote N–dimensional vector of the noise samples $e = [e_1 \quad \ldots e_N]^\top$. Thinking of f as an infinite–dimensional column vector, and using notation of ordinary algebra to handle infinite–dimensional objects, model (1) can be expressed in matrix form as

$$y = Gf + e \tag{2}$$

where $G \in \mathbb{R}^{N \times \infty}$ is a matrix whose entries are defined by the system input, so that Gf represents the convolution between the system impulse response and the input. We consider the problem of estimating f from y.

2.2 Gaussian process regression via Stable Spline Kernels

In the classical system identification set up, the impulse response is searched for within a finite–dimensional space, e.g. postulating ARX, ARMAX or Laguerre models. Under the framework of Gaussian process regression [Rasmussen and Williams, 2006], f is instead modeled as a sampled version of a continuous–time zero–mean Gaussian process with a suitable covariance (kernel), independent of e. We denote with K the infinite–dimensional matrix obtained by sampling $K(\cdot, \cdot)$ on $\mathbb{N} \times \mathbb{N}$ and write

$$f \sim \mathcal{N}(0, K(\eta)), \qquad f \perp e \tag{3}$$

where η is a vector of hyperparameters governing the prior covariance. According to an Empirical Bayes paradigm [Berger, 1985, Maritz and Lwin, 1989], the hyperparameters can be estimated from the data via marginal likelihood maximization, i.e. by maximizing the marginalization with respect to f of the joint density of y and f

$$\hat{\eta} = \arg\min_{\eta} \left\{ \log\det \Sigma_y(\eta) + y^\top \Sigma_y(\eta)^{-1} y \right\} \tag{4}$$

with

$$\Sigma_y(\eta) = GK(\eta)G^\top + \sigma^2 I_N \,. \tag{5}$$

Once η is estimated, the impulse response can be computed as the minimum variance estimate given y and $\hat{\eta}$, i.e.

$$\hat{f} := \mathbb{E}\left[f|y,\hat{\eta}\right] = K(\hat{\eta})G^{\top}\left(GK(\hat{\eta})G^{\top} + \sigma^2 I_N\right)^{-1} y\,. \tag{6}$$

Prior information is introduced in the identification process by assigning the covariance $K(\eta)$. The quality of the estimates crucially depends on this choice as well as on the quality of the estimated $\hat{\eta}$.

A class of prior covariances which has been proved to be very effective in the system identification scenario, is the class of stable spline kernels ([Pillonetto and De Nicolao, 2010, Pillonetto et al., 2010, 2011]), that, besides incorporating information on smoothness, guarantees that the estimated impulse response is almost surely stable.

First–order stable spline kernels (equivalently, stable spline kernels of order 1) were introduced in [Pillonetto et al., 2010] (see also [Chen et al., 2012], where they are referred to as Tuned/Correlated (TC) kernels) and are defined as

$$\mathcal{K}_{ij} = \lambda\alpha^{\max(i,j)}, \qquad \lambda \geq 0,\ 0 \leq \alpha < 1,$$

so that $\eta = [\lambda, \alpha]$.

3 Maximum Entropy band extension problem

Covariance extension problems were introduced by A. P. Dempster [Dempster, 1972] and studied by many authors (see e.g. [Grone et al., 1984, Dym and Gohberg, 1981, Johnson, 1990, Gohberg et al., 1993, Dahl et al., 2008] and references therein, see also [Carli et al., 2011, Carli and Georgiou, 2011, Carli et al., 2013] for an extension to the circulant case). In the literature concerning matrix completion problems, it is common practice to describe the pattern of the specified entries of an $n \times n$ partial symmetric matrix by an undirected graph of n vertices which has an edge joining vertex i and vertex j if and only if the (i, j) entry is specified. If the graph of the specified entries is *chordal* (i.e., a graph in which every cycle of length greater than three has an edge connecting nonconsecutive nodes, see e.g. [Golumbic, 1980]), and, in particular, if the specified elements lie on a band centered along the main diagonal, then the maximum entropy covariance extension problem admits a closed form solution in terms of the principal minors of the matrix to be completed (see [Barrett et al., 1989], [Fukuda et al., 2000], [Nakata et al., 2003]). In this section, we briefly review some fundamental results about *maximum entropy band extension problems* that will be used to prove our main results in Section 4.

Recall that the *differential entropy* $H(p)$ of a probability density function p on $\mathbb{R}^n$ is defined by

$$H(p) = -\int_{\mathbb{R}^n} \log(p(x))p(x)dx. \tag{7}$$

In case of a zero–mean Gaussian distribution p with covariance matrix $\mathbf{\Sigma}_n$, we get

$$H(p) = \frac{1}{2}\log(\det \mathbf{\Sigma}_n) + \frac{1}{2}n\left(1 + \log(2\pi)\right). \tag{8}$$

Let $\mathcal{I} \subset \{1,\dots,n\} \times \{1,\dots,n\}$ denote a set of indices and $\bar{\mathcal{I}}$ the complement of $\mathcal{I}$ with respect to $\{1,\dots,n\} \times \{1,\dots,n\}$. Let x be the vector, say k–dimensional, obtained by

stacking the x_{ij}'s one on top of the other. A *partial matrix* is a parametric family of $n \times n$ matrices $\mathbf{\Sigma}_n(x)$ with entries $[\mathbf{\Sigma}_n(x)]_{i,j} = \sigma_{ij}$, $(i,j) \in \mathcal{I}$ specified, and entries $[\mathbf{\Sigma}_n(x)]_{i,j} = x_{ij}$, for $(i,j) \in \bar{\mathcal{I}}$, which are left unspecified. Here, both σ_{ij} and x_{ij} are taken to be real. A *completion (extension)* of $\mathbf{\Sigma}_n(x)$ is a $n \times n$ matrix $[C]_{i,j} = c_{ij}$ which satisfies

$$c_{ij} = \sigma_{ij} \quad \forall (i,j) \in \mathcal{I}\,.$$

In particular, let

$$\mathcal{I}_b^{(m)} := \{(i,j) \mid |i-j| \leq m\}\,.$$

If $\mathcal{I} \equiv \mathcal{I}_b^{(m)}$, we refer to $\mathbf{\Sigma}_n^{(m)}(x)$ as a *partially specified m–band matrix*.

Consider the following optimization problem

$$\text{minimize} \quad \left\{ -\log\det \mathbf{\Sigma}_n^{(m)}(x) \mid \mathbf{\Sigma}_n^{(m)}(x) \in \mathcal{S}_n \right\} \tag{9a}$$

$$\text{subject to} \quad \mathbf{\Sigma}_n^{(m)}(x) \succeq 0 \tag{9b}$$

$$e_i^\top\, \mathbf{\Sigma}_n^{(m)}(x)\, e_j = \sigma_{ij}, \quad (i,j) \in \mathcal{I}_b^{(m)} \tag{9c}$$

with optimization variable x, namely the problem of computing the maximum entropy extension of the partially specified symmetric m–band matrix $\mathbf{\Sigma}_n^{(m)}(x)$. Problem (9) is a convex optimization problem. Denote by x^o its optimal value and by $\mathbf{\Sigma}_n^{(m),o} \equiv \mathbf{\Sigma}_n^{(m)}(x^o)$ the associated extension. Moreover from now on, we will drop the dependence on x in $\mathbf{\Sigma}_n^{(m)}(x)$ and refer to a m-band partially specified $n \times n$ matrix as $\mathbf{\Sigma}_n^{(m)}$.

Theorem 3.1 ([Dempster, 1972, Dym and Gohberg, 1981])**.** (*i*) Feasibility: *Problem* (9) *is feasible, namely $\mathbf{\Sigma}_n^{(m)}$ admits a positive definite extension if and only if*

$$\begin{bmatrix} \sigma_{i,i} & \cdots & \sigma_{i,m+i} \\ \vdots & & \vdots \\ \sigma_{m+i,i} & \cdots & \sigma_{m+i,m+i} \end{bmatrix} \succ 0, \quad i = 1, \ldots, n-m \tag{10}$$

(*ii*) Bandedness: *Assume* (10) *holds. Then* (9) *admits a unique solution with the additional property that its inverse is banded of bandwidth m, namely the (i,j)–th entry of $\left(\mathbf{\Sigma}_n^{(m),o}\right)^{-1}$ is zeros if $|i-j| > m$.*

The positive definite maximum entropy extension $\mathbf{\Sigma}_n^{(m),o}$ is also called *central extension* of $\mathbf{\Sigma}_n^{(m)}$.

Let $\bar{\Sigma}$ be such that $\left[\bar{\Sigma}\right]_{ij} = \sigma_{ij} \quad (i,j) \in \mathcal{I}_b^{(m)}$. Then, it can be shown [Dempster, 1972, Dahl et al., 2008] that Problem (9) is equivalent to the following optimization problem

$$\text{minimize} \quad \log\det \mathbf{\Sigma}_n^{(m)} + \operatorname{trace}\left(\bar{\Sigma}\left(\mathbf{\Sigma}_n^{(m)}\right)^{-1}\right) \tag{11a}$$

$$\text{subject to} \quad \mathbf{\Sigma}_n^{(m)} \succeq 0 \tag{11b}$$

$$e_i^\top \left(\mathbf{\Sigma}_n^{(m)}\right)^{-1} e_j = 0, \quad (i,j) \in \bar{\mathcal{I}}_b^{(m)} \tag{11c}$$

If we denote with $\theta = \left[\theta_1, \ldots, \theta_n\right]^\top$ a zero–mean Gaussian random vector with covariance $\mathbf{\Sigma}_n^{(m)}$, then (11c) holds if and only if the random variables θ_i, θ_j in θ are conditionally independent given the others (see e.g. [Dempster, 1972]). In other words, if we denote with $\bar{\mathbf{\Sigma}}$ the sample covariance of θ, the equivalence between Problem (9) and Problem (11) states that the covariance matrix that maximizes the entropy among all the covariance matrices with given first $m+1$ covariance lags, is also the one that maximizes the likelihood among all the covariance matrices satisfying the conditional independence constraints (11c).

For banded sparsity pattern like those considered so far, Problem (9) admits a closed form solution that can be computed recursively in the following way. We start by considering a partially specified $n \times n$ symmetric matrix of bandwidth $(n-2)$

$$\mathbf{\Sigma}_n^{(n-2)} = \begin{bmatrix} \sigma_{1,1} & \sigma_{1,2} & \cdots & \sigma_{1,n-1} & x \\ \sigma_{1,2} & \sigma_{2,2} & \cdots & \sigma_{2,n-1} & \sigma_{2,n} \\ \vdots & \vdots & & \vdots & \vdots \\ \sigma_{1,n-1} & \sigma_{2,n-1} & \cdots & \sigma_{n-1,n-1} & \sigma_{n-1,n} \\ x & \sigma_{2,n} & \cdots & \sigma_{n-1,n} & \sigma_{n,n} \end{bmatrix} \tag{12}$$

and consider the submatrix

$$L = [\sigma_{ij}]_{i,j=1}^{n-1}. \tag{13}$$

We call *one–step extensions* the extensions of $n \times n$, $(n-2)$–band matrices. The following theorem gives a recursive algorithm to compute the extension of partially specified matrices of generic bandwidth m by computing the one–step extensions of suitable submatrices. It also gives a representation of the solution in factored form.

Theorem 3.2 ([Gohberg et al., 1993], [Dym and Gohberg, 1981])**.** *(i) The one–step central extension of $\mathbf{\Sigma}_n^{(n-2)}$ is given by*

$$x^o = -\frac{1}{y_1} \sum_{j=2}^{n-1} \sigma_{nj} y_j \tag{14}$$

with

$$\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_{n-1} \end{bmatrix} = L^{-1} \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \tag{15}$$

Let $\mathbf{\Sigma}_n^{(m)}$ be an $n \times n$ partially specified m–band matrix. The central extension $C = [c_{ij}]_{i,j=1}^n$ of $\mathbf{\Sigma}_n^{(m)}$ is such that for all $m+1 < t \leq n$ and $1 \leq s \leq t-m-1$ the submatrices

$$C(\{s, \ldots, t\}) = \begin{bmatrix} c_{s,s} & \cdots & c_{s,t} \\ \vdots & & \vdots \\ c_{t,s} & \cdots & c_{t,t} \end{bmatrix} \tag{16}$$

are the central one–step extensions of the corresponding $(t-s-1)$–band matrix.

(*ii*) *In particular, the central extension of the partially specified symmetric m-band matrix* $\mathbf{\Sigma}_n^{(m)}$ *admits the factorization*

$$\mathbf{C} = \left(L_n^{(m)} V U_n^{(m)} \right)^{-1} \tag{17}$$

where $L_n^{(m)} = [\ell_{ij}]$ *is a lower triangular banded matrix with ones on the main diagonal,* $\ell_{jj} = 1$*, for* $j = 1, \dots, n$*, and*

$$\begin{bmatrix} \ell_{\alpha j} \\ \vdots \\ \ell_{\beta j} \end{bmatrix} = - \begin{bmatrix} \sigma_{\alpha\alpha} & \dots & \sigma_{\alpha\beta} \\ \vdots & \ddots & \vdots \\ \sigma_{\beta\alpha} & \dots & \sigma_{\beta\beta} \end{bmatrix}^{-1} \begin{bmatrix} \sigma_{\alpha j} \\ \vdots \\ \sigma_{\beta j} \end{bmatrix} \tag{18}$$

for $j = 1, \dots, n-1$*,* $U_n^{(m)} = \left(L_n^{(m)} \right)^\top$ *and* $V = [v_{ij}]$ *diagonal with entries*

$$v_{jj} = \left(\begin{bmatrix} \sigma_{jj} & \dots & \sigma_{j\beta} \\ \vdots & \ddots & \vdots \\ \sigma_{\beta j} & \dots & \sigma_{\beta\beta} \end{bmatrix}^{-1} \right)_{1,1} , \tag{19}$$

for $j = 1, \dots, n$*, where*

$$\begin{array}{ll} \alpha = \alpha(j) = j + 1 & \text{for} \quad j = 1, \dots, n-1 \,, \\ \beta = \beta(j) = \min(j + m, n) & \text{for} \quad j = 1, \dots, n \,. \end{array}$$

4 Maximum Entropy properties of the First-order Stable Spline kernel and Its Implications

In this section, we provide an independent proof of the maximum entropy property of first–order stable spline kernels that relies on the theory of matrix extension problems introduced in the previous section. This argument leads to a closed form expression for the inverse of the first order stable spline kernel as well as to a new factorization. Maximum likelihood properties of the stable spline kernel are also highlighted.

Proposition 4.1. *Consider Problem* (9) *with* $m = 1$ *and*

$$\sigma_{ij} = \mathcal{K}_{ij} = \alpha^{\max(i,j)}, \quad (i,j) \in \mathcal{I}_b^{(1)} \tag{20}$$

i.e. consider the partially specified 1*–band matrix*

$$\mathbf{\Sigma}_n^{(1)}(x) = \begin{bmatrix} \alpha & \alpha^2 & x_{13} & \dots & \dots & x_{1n} \\ \alpha^2 & \alpha^2 & \alpha^3 & x_{24} & \dots & x_{2n} \\ x_{13} & \alpha^3 & \alpha^3 & \alpha^4 & & \vdots \\ \vdots & & \ddots & \ddots & \ddots & x_{n-2,1} \\ \vdots & & & \alpha^{n-1} & \alpha^{n-1} & \alpha^n \\ x_{1n} & \dots & \dots & x_{n-2,1} & \alpha^n & \alpha^n \end{bmatrix}$$

Then $\mathbf{\Sigma}_n^{(1)}(x^o) = \mathcal{K}$*, i.e. the solution of the Maximum Entropy Problem* (9) *coincides with the first order stable spline kernel.*

Proof. By Theorem 3.2, the maximum entropy completion of $\mathbf{\Sigma}_n^{(1)}(x)$ can be recursively computed starting from the maximum entropy completions of the nested principal submatrices of smaller size. The statement can thus be proved by induction on the dimension n of the completion.

- Let $n = 3$, then by (14)–(15), the central extension of

$$\begin{bmatrix} \alpha & \alpha^2 & x_{13} \\ \alpha^2 & \alpha^2 & \alpha^3 \\ x_{13} & \alpha^3 & \alpha^3 \end{bmatrix}$$

  is given by $x_{13}^o = \alpha^3 = \mathcal{K}(1,3)$, as claimed.

- Now assume that the statement holds for $n = k$, $k \geq 3$, i.e. that $\mathcal{K}(\{1,\dots,k\})$ is the central extension of $\mathbf{\Sigma}_n^{(1)}(\{1,\dots,k\})$. We want to prove that $\mathcal{K}(\{1,\dots,k+1\})$ is the central extension of $\mathbf{\Sigma}_n^{(1)}(\{1,\dots,k+1\})$. To this aim, we only need to prove that the $(k-1)$ submatrices $\mathcal{K}(\{s,\dots,k+1\})$, $1 \leq s \leq k-1$, are the central one–step extensions of the corresponding $(k-s)$–band matrices

$$\begin{bmatrix} \alpha^s & \dots & \alpha^k & x_{s,k+1} \\ \vdots & \ddots & & \alpha^{k+1} \\ & & \ddots & \vdots \\ x_{s,k+1} & \alpha^{k+1} & \dots & \alpha^{k+1} \end{bmatrix}.$$

  or, equivalently, that $x_{s,k+1}^o = \mathcal{K}(s,k+1)$, for $s = 1,\dots,k-1$. In order to find $x_{s,k+1}^o$, we consider (15), which, by the inductive hypothesis, becomes

$$\begin{bmatrix} y_1^{(s,k+1)} \\ y_2^{(s,k+1)} \\ \vdots \\ y_k^{(s,k+1)} \end{bmatrix} = \mathcal{K}(\{s,\dots,k\})^{-1} \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

  By considering the adjoint of $\mathcal{K}(\{s,\dots,k\})$ one can see that $y_2^{(s,k+1)} = -y_1^{(s,k+1)}$ while all the others $y_i^{(s,k+1)}$, $i = 3,\dots,k$ are identically zero. It follows that

$$x_{s,k+1}^o = -\frac{1}{y_1^{(s,k+1)}} y_2^{(s,k+1)} \alpha^{k+1} = \alpha^{k+1} = \mathcal{K}(s,k+1),$$

  as claimed.

□

From the equivalence between the maximum entropy problem 9 and the maximum likelihood problem 11 we get the following.

Proposition 4.2. *Let $m = 1$ and $\left[\bar{\Sigma}\right]_{ij} = \left[\theta\theta^\top\right]_{ij} = \alpha^{\max(i,j)}$, $(i,j) \in \mathcal{I}_b^{(1)}$, then the first-order stable spline kernel maximizes the likelihood in* (11a) *among all covariances that satisfies* (11b) *and the conditional independence constraints* (11c).

Proposition 4.3. *The following are equivalent*

(i) $\mathcal{K}$ *solves Problem* (9) *with* $m = 1$ *and moment constraints as in* (20).

(ii) $\mathcal{K}$ *admits the factorization*

$$\mathcal{K} = UWU^\top \tag{21}$$

with

$$U = \begin{bmatrix} 1 & 1 & \dots & 1 \\ 0 & 1 & \dots & 1 \\ 0 & 0 & \ddots & \vdots \\ 0 & \dots & 0 & 1 \end{bmatrix}, \tag{22}$$

and

$$W = (\alpha - \alpha^2)\text{diag}\left\{1, \alpha, \alpha^2, \dots, \alpha^{n-2}, \frac{\alpha^{n-1}}{1-\alpha}\right\}. \tag{23}$$

(iii) $\mathcal{K}^{-1}$ *is tridiagonal banded and is given by*

$$\frac{1}{\alpha - \alpha^2}\begin{bmatrix} 1 & -1 & 0 & \dots & 0 \\ -1 & 1 + \frac{1}{\alpha} & -\frac{1}{\alpha} & \ddots & \vdots \\ 0 & & \frac{1}{\alpha} + \frac{1}{\alpha^2} & & 0 \\ \vdots & \ddots & & \ddots & -\frac{1}{\alpha^{n-2}} \\ 0 & \dots & 0 & -\frac{1}{\alpha^{n-2}} & \frac{1}{\alpha^{n-2}} + \frac{1-\alpha}{\alpha^{n-1}} \end{bmatrix} \tag{24}$$

Proof. That $\mathcal{K}$ admits the factorization (21)–(23) follows from Theorem 3.2 (ii). In fact, by (17)–(19) the inverse of the stable spline kernel of order 1 can be factored as

$$\mathcal{K}^{-1} = L_n^{(1)} V U_n^{(1)} \tag{25}$$

where $L_n^{(1)}$ takes the form

$$L_n^{(1)} = \left(U_n^{(1)}\right)^\top = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ -1 & 1 & 0 & \dots & 0 \\ 0 & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix} \tag{26}$$

and

$$V = \frac{1}{\alpha - \alpha^2}\text{diag}\left\{1, \frac{1}{\alpha}, \frac{1}{\alpha^2}, \dots, \frac{1}{\alpha^{n-2}}, \frac{1-\alpha}{\alpha^{n-1}}\right\} \tag{27}$$

Bandedness of $\mathcal{K}^{-1}$ follows from Theorem 3.1 (ii) and expression (24) for $\mathcal{K}^{-1}$ is an immediate consequence of the factorization (21)–(23). □

Remark 4.1. *Let* $A = \begin{bmatrix} I_{n-1} & | & 0_{n-1} \end{bmatrix}$. *It is an immediate consequence of* (24) *that*

$$\left(A\mathcal{K}^{-1}\right)\mathbb{1}_{n-1} = 0_{n-1} \tag{28}$$

i.e. the first $n-1$ *columns of* $\mathcal{K}^{-1}$ *sum up to zero.*

Corollary 4.1. *The stable spline kernel of order* 1 *has determinant*

$$\det(\mathcal{K}) = \left[(1-\alpha)^{n-1}\alpha^{\frac{1}{2}n(n+1)}\right]. \tag{29}$$

Proof. The first (resp., third) factor in the right hand side of (25) is a lower (resp., upper) triangular matrix with diagonal entries equal to one, and hence the positive definite matrix $\mathcal{K}^{-1}$ and V have the same determinant, i.e.

$$\det(\mathcal{K}^{-1}) = \det(V) = \frac{1}{\alpha^n(\alpha-\alpha^2)^{n-1}}\prod_{i=2}^{n}\frac{1}{\alpha^{i-2}}.$$

The thesis follows immediately by recalling that $\sum_{i=1}^{n-2} i = \frac{1}{2}(n-2)(n-1)$. □

Remark 4.2. *A key point in the evaluation of the stable spline estimator lies in solving the marginal likelihood maximization problem* (4)*, that is usually nonconvex. No matter what solver is used, the tuning of the hyperparameters requires repeated evaluations of the marginal likelihood. Here we observe that, whatever the value of* α*, all the stable spline kernels of order* 1 *share the same factor* U (22)*. Moreover, being* W *available in closed form, once* α *is known the factorization* (21) *is computationally inexpensive. The same applies to the factorization of* $\mathcal{K}^{-1}$*. This fact, together with the closed form expression for the determinant of the stable spline kernel in* (29)*, can be exploited both to improve the stability and to reduce the computational burden associated with computational schemes for the evaluation of the stable spline estimator like those in [Carli et al., 2012b, Chen and Ljung, 2013].*

We conclude this section by highlighting an additional property of the first–order stable spline kernel that originates from the maximum entropy property of Proposition 4.1.

5 Conclusions

Empirical Bayes estimation for system identification problems has recently become popular, mainly due to the introduction of a family of prior descriptions (the so–called stable spline kernels) which encode structural properties of dynamical systems such as stability. Maximum entropy properties of first–order stable spline kernels have been highlighted in [Pillonetto and De Nicolao, 2011]. In this paper we provide an alternative proof that leads to a closed form expression for the inverse of the first order stable spline kernel as well as to a new, computationally advantageous factorization. Maximum likelihood properties of the stable spline kernel are also highlighted. These properties can be exploited both to improve the stability and to relieve the computational complexity associated with the computation of stable spline estimators.

References

H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19:716–723, 1974.

A. Aravkin, J. V. Burke, A. Chiuso, and G. Pillonetto. Convex vs non-convex estimators for regression and sparse estimation: the mean squared error properties of ard and glasso. *Journal of Machine Learning Research*, 15:217–252, 2014.

W.W. Barrett, C.R. Johnson, and M. Lundquist. Determinantal formulation for matrix completions associated with chordal graphs. *Linear Algebra and its Applications*, 121: 265–289, 1989.

J. O. Berger. *Statistical decision theory and Bayesian analysis*. Springer Verlag, 1985.

F. P. Carli, T. Chen, A. Chiuso, L. Ljung, and G. Pillonetto. On the estimation of hyperparameters for bayesian system identification with exponentially decaying kernels. In *Proceedings of the 51st IEEE Conference on Decision and Control (CDC 2012)*, pages 5260–5265. IEEE, 2012a.

F. P. Carli, A. Chiuso, and G. Pillonetto. Efficient algorithms for large scale linear system identification using stable spline estimators. In *Proceedings of the 16th IFAC symposium on system identification (SYSID 2012)*, pages 119–124. IFAC, 2012b.

F.P. Carli and T.T. Georgiou. On the covariance completion problem under a circulant structure. *IEEE Transactions on Automatic Control*, 56(4):918 – 922, 2011.

F.P. Carli, A. Ferrante, M. Pavon, and G. Picci. A maximum entropy solution of the covariance extension problem for reciprocal processes. *IEEE Transactions on Automatic Control*, 56(9):1999–2012, 2011.

F.P. Carli, A. Ferrante, M. Pavon, and G. Picci. An efficient algorithm for maximum entropy extension of block-circulant covariance matrices. *Linear Algebra and its Applications*, 439 (8):2309–2329, 2013.

T. Chen and L. Ljung. Implementation of algorithms for tuning parameters in regularized least squares problems in system identification. *Automatica*, 49(7):2213–2220, 2013.

T. Chen, H. Ohlsson, G. C. Goodwin, and L. Ljung. Kernel selection in linear system identification part II: A classical perspective. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC 2011)*, pages 4326–4331. IEEE, 2011.

T. Chen, H. Ohlsson, and L. Ljung. On the estimation of transfer functions, regularizations and Gaussian processes - revisited. *Automatica*, 48(8):1525–1535, 2012.

J. Dahl, L. Vanderberghe, and V. Roychowdhury. Covariance selection for non–chordal graphs via chordal embedding. *Optimization Methods and Software*, 23:501–520, 2008.

A. P. Dempster. Covariance selection. *Biometrics*, 28:157–175, 1972.

H. Dym and I. Gohberg. Extensions of band matrices with band inverses. *Linear algebra and its applications*, 36:1–24, 1981.

M. Fukuda, M. Kojima, K. Murota, and K. Nakata. Exploiting sparsity in semidefinite programming via matrix completion i: general framework. *SIAM Journal on Optimization*, 11:647–674, 2000.

I. Gohberg, S. Goldberg, and M. A. Kaashoek. *Classes of Linear Operators, II,*. Birkhäuser, Basel, 1993.

M. Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Academic Press, New York, 1980.

R. Grone, C.R. Johnson, E.M. Sa, and H. Wolkowicz. Positive definite completions of partial Hermitian matrices. *Linear Algebra and Its Applications*, 58:109–124, 1984.

T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning*. Springer, 2008.

C.R. Johnson. Matrix completion problems: a survey. In *Proceedings of Symposia in Applied Mathematics (1990)*, volume 40, pages 171–198, 1990.

L. Ljung. *System Identification - Theory For the User*. Prentice Hall, 1999.

J. S. Maritz and T. Lwin. *Empirical Bayes methods*. Chapman and Hall London, 1989.

K. Nakata, K. Fujitsawa, M. Fukuda, M. Kojima, and K. Murota. Exploiting sparsity in semidefinite programming via matrix completion ii: implementation and numerical details. *Mathematical Programming Series B*, 95:303–327, 2003.

G. Pillonetto and G. De Nicolao. A new kernel-based approach for linear system identification. *Automatica*, 46:81–93, 2010.

G. Pillonetto and G. De Nicolao. Kernel selection in linear system identification Part I: A Gaussian process perspective. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 4318–4325. IEEE, 2011.

G. Pillonetto, A. Chiuso, and G. De Nicolao. Regularized estimation of sums of exponentials in spaces generated by stable spline kernels. In *American Control Conference (ACC), 2010*, pages 498–503. IEEE, 2010.

G. Pillonetto, A. Chiuso, and G. De Nicolao. Prediction error identification of linear systems: A nonparametric Gaussian regression approach. *Automatica*, 47:291–305, 2011.

C. Rasmussen and C. Williams. *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.

B. Scholkopf and A. J. Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.

G. Schwarz. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.

J. Shawe-Taylor and N. Cristianini. *Kernel methods for pattern analysis*. Cambridge University Press, 2004.

T. Soderstrom and P. Stoica. *System Identification*. Prentice Hall, 1989.